arXiv:1612.04898v1 [stat.ML] 15 Dec 2016

Efficient Distributed Semi-Supervised Learning using Stochastic Regularization over Affinity Graphs

Sunil Thulasidasan[1,2]
sunil@lanl.gov

Jeffrey Bilmes[2]
bilmes@uw.edu

Garrett Kenyon[1]
gkenyon@lanl.gov

[1]Los Alamos National Laboratory
[2]Department of Electrical Engineering, University of Washington

Abstract

We describe a computationally efficient, stochastic graph-regularization technique that can be utilized for the semi-supervised training of deep neural networks in a parallel or distributed setting. We utilize a technique, first described in [13] for the construction of mini-batches for stochastic gradient descent (SGD) based on synthesized partitions of an affinity graph that are consistent with the graph structure, but also preserve enough stochasticity for convergence of SGD to good local minima. We show how our technique allows a graph-based semi-supervised loss function to be decomposed into a sum over objectives, facilitating data parallelism for scalable training of machine learning models. Empirical results indicate that our method significantly improves classification accuracy compared to the fully-supervised case when the fraction of labeled data is low, and in the parallel case, achieves significant speed-up in terms of wall-clock time to convergence. We show the results for both sequential and distributed-memory semi-supervised DNN training on a speech corpus.

1 Introduction

Big data is often a deluge of unstructured, un-annotated and unlabeled data - a natural outcome of technological advances that have enabled data to be collected and disseminated with little effort and on large scales but where annotating and labeling ground truth largely remains a time consuming, human effort. Semi-supervised learning (SSL) methods use both labeled and unlabeled data to improve learning performance [2] and are especially useful in situations where labeled data is scarce Such methods leverage unlabeled data by exploiting the similarity between labeled and unlabeled data by capturing this relationship via graphs, where the nodes represent both labeled and unlabeled points and the weights of the edges reflect the similarity between the nodes [16].

The main idea behind graph-based SSL methods is that given a similarity metric, the objective function constrains similar (i.e., nearby) nodes to have the same label by imposing a graph-neighbor regularization. This is effective because it forces the labels to be consistent with the graph structure, and the underlying manifold represented thereby. The general form of the loss function in graph-based SSL has the following form

$$\sum_{i=1}^{L} l(y_i, f(x_i)) + \lambda \sum_{i,j} \omega_{i,j} g(f(x_i), f(x_j)) \tag{1}$$

where $f : \mathcal{X} \rightarrow \mathcal{Y}$ is the classifier mapping from input to output space. The first term in Equation 1 the supervised loss function calculated on the labeled points, which can be a squared loss, hinge-

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

loss or some measure of divergence between predictions and ground truth. The second term is the graph regularizer, where $\omega_{i,j}$ captures the similarity between points x_i and x_j. $g(.)$ captures the discrepancy between output $f(x_i)$ and $f(x_j)$, incurring a penalty when similar nodes have differing outputs. Additional regularizers such as the standard ℓ_2 regularizer can also be applied to the above loss function to prevent overfitting. Concretely, Let $\{(\mathbf{x_i}, \mathbf{y_i})\}_{\mathbf{i=1}}^{\ell}$ be the labeled training data and $\{\mathbf{x_i}\}_{\mathbf{i=\ell+1}}^{\ell+\mathbf{u}}$ be the unlabeled training data, where $n = \ell + u$ so that we have n points in total. We assume that the samples $\{\mathbf{x}_i\}_i$ are used to produce a weighted undirected graph $\mathcal{G} = (V, E, \mathbf{W})$, where $\omega_{i,j} \in \mathbf{W}$ is taken to be the similarity (edge weight) between samples (vertices) $\mathbf{x}_i$ and $\mathbf{x}_j$. We use the objective function defined in [12, 10], namely:

$$J(\theta) = \sum_{i=1}^{l} \mathbf{D}(\mathbf{t}_i \parallel \mathbf{p}_\theta(\mathbf{x}_i)) + \gamma \sum_{i,j=1}^{n} \omega_{i,j} \mathbf{D}(\mathbf{p}_\theta(\mathbf{x}_i) \parallel \mathbf{p}_\theta(\mathbf{x}_j)) \quad + \kappa \sum_{i=1}^{n} \mathbf{D}(\mathbf{p}_\theta(\mathbf{x}_i) \parallel \mathbf{u}) + \lambda \|\theta\|, \tag{2}$$

where $J(\theta)$ is the loss calculated over all samples. We use KL-divergence (denoted by $\mathbf{D}(. \parallel .)$) as our loss function since our output is a probability distribution over classes. The first term in the above equation is the supervised loss over the training samples, and the second term is the penalty imposed by the graph regularizer over neighboring pairs of nodes that favors smooth solutions over the graph. The third term is an entropy regularizer ($\mathbf{u}$ is the uniform distribution) and favors higher entropy distributions to discourage degenerate solutions. The final term in Equation 2 is the standard ℓ_2 regularizer to discourage overfitting. Note that the loss function, as such, is not directly decomposable as a sum over data points due to the presence of the graph regularizer and thus is not directly amenable to data parallelism. This necessitates the implementation of data partitioning; such strategies for parallel machine learning in the fully supervised case have been described, for example, in [8, 6]. The presence of the graph regularizer term in the semi-supervised case necessitates a different approach, and in this work we study the effectiveness of a stochastic method in the parallel setting, first described in [13]. Our method constructs graph-based mini-batches by sampling the data using graph partitioning, but at the same time also preserves the statistical properties of the data distribution. For the experiments in this work, we use a deep neural network with a loss function given by the above equation but the method can be generalized to any parametric learner.

1.1 Graph Partitioning for Objective Function Decomposition

Decomposing KL-divergence in Equation 2 into entropy and cross-entropy terms and dropping the constant terms (w.r.t parameters), we can show that over one labeled point, the loss function becomes

$$J_i = H^c(\mathbf{t}_i, \mathbf{p}_i) + \gamma \sum_{j=1}^{n} \omega_{ij} H^c(\mathbf{p}_i, \mathbf{p}_j) - (\kappa + \gamma \sum_{j=1}^{n} \omega_{ij}) H(\mathbf{p}_i) + \lambda \|\theta\| \tag{3}$$

where H and H^c are, respectively, the entropy and cross-entropy. Since we are dealing with a non-convex objective function, and a moderately large data set ($\approx$ 1 million training samples), we use stochastic gradient descent to optimize our objective function. We also use mini-batches to improve the gradient quality, and further, use larger mini-batches (size set either to 1024 or 2048) for better computational efficiency on GPUs. In order to converge to good local minima, traditional SGD methods require randomly shuffling the data before constructing the mini-batches; this, however, poses a serious problem for our objective function. To see this, consider the terms involving graph regularization from our decomposed objective function, calculated over each point:

$$G_i = \gamma \sum_{j=1}^{n} \omega_{ij} H^c(\mathbf{p}_i, \mathbf{p}_j) - \gamma \sum_{j=1}^{n} \omega_{ij} H(\mathbf{p}_i) \tag{4}$$

For the graph regularization term to have any effect at all, the w_{ij}'s corresponding to the points in the mini-batch have to be non-zero. For a randomly shuffled data-set, given that the graph is very sparse (since each of the $\approx$ 1 million points only has a little more than 10 neighbors), the chunk of the affinity matrix corresponding to the mini-batch will be extremely sparse, implying that graph regularization will fail to take place on most computations. One way to fix this is, for a given mini-batch, to loop over all the neighbors for each point in the mini-batch, but this prevents us from doing efficient matrix-matrix multiplications and completely degrades performance negating any benefits of using

fast processors like GPUs. Thus, for the graph regularizer to be effective in a computationally efficient way, our mini-batches need to reflect the structure of the graph. To do this, we partition our affinity graph into k balanced parts (by minimizing edge-cut) which results in a re-permuted affinity matrix that has a dense block-diagonal structure as shown in Figure 1b; contrast this with the affinity matrix before partitioning in Figure 1a where most entries over a 1000×1000 block (corresponding to a mini-batch size of 1000) are zero.

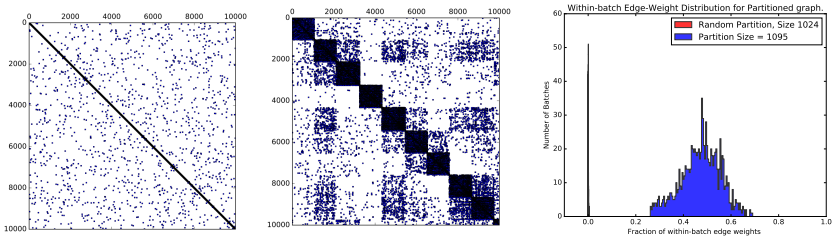
Within-batch Edge-Weight Distribution for Partitioned graph.
Random Partition, Size 1024
Partition Size = 1095
Number of Batches
Fraction of within-batch edge weights

(a) (b) (c)

Figure 1: (1a) $10,000 \times 10,000$ sub-block of the affinity matrix for randomly shuffled data. Choosing 1000×1000 blocks along the diagonal of such a matrix would produce extremely sparse matrices. (1b) Same affinity matrix re-permuted based on graph partitioning. While perform mini-batch computation we choose the diagonal blocks (dense squares). (1c) Distribution showing the within-batch connectivity of graph-based batches (blue) vs randomly shuffled batches (spike). Nodes in randomly shuffled batches have almost none of the neighbors in the same batch

Dense mini-batches imply that most of the neighbors of the nodes within a mini-batch are contained within the same batch. More formally, let $\mathcal{N}_i$ represent the set of neighbors of node i and $\mathcal{C}_i \subseteq \mathcal{N}_i$ be the set of neighbors of a node i that are within the same batch. Let $\mathcal{M}_j$ be the set that represents mini-batch j. We define the within-batch connectivity of $\mathcal{M}_j$ as

$$c_j = \frac{\sum |\mathcal{C}_i|}{\sum |\mathcal{N}_i|}, \forall i \in \mathcal{M}_j, j = 1, 2, 3 \ldots k \tag{5}$$

In the randomly shuffled (pre-partitioning scenario) we expect most of the c_j's to be close to zero, while for graph partitioned mini-batches, we expect a relatively higher c. Figure 1c shows this distribution for the random minibatches (seen as a sharp spike near 0) and for graph-partitioned mini-batches (in blue). Partitioning gives us an efficient way of computing our objective function: given a matrix permutation induced by the graph partitioning, we re-permute the affinity and data matrices accordingly. Then, during each mini-batch computation, we calculate the objective function and gradients on these partitions. Note that graph-partitioning is a pre-processing operation, and only done once before training commences.

2 Issues for Stochastic Optimization

Theoretically, SGD, gives us an unbiased estimate of the true gradient, but only if the data is sampled form the true distribution. If our entire data set approximates this true distribution reasonably well, then a randomly sampled mini-batch will also be faithful to this distribution. However, for a graph partitioned mini-batch this argument no longer holds since the data points that comprise a mini-batch are not randomly sampled, but on the contrary, reflect relatively homogeneous regions on some low dimensional manifold (since we are partitioning by minimizing edge-cut on a k-NN graph). Thus our gradient estimate is no longer unbiased, leading to poor convergence of SGD. On the other hand we have also seen that randomly shuffled batches will cause the graph regularizer to become ineffective due to poor within-batch neighbor connectivity, unless one accepts extremely long computational times (and communication costs in a parallel implementation). For a more formal discussion see [13].

2.1 Improving SGD Convergence using Graph-Synthesized Meta-batches

Constructing a mini-batch that gives good SGD convergence, and good neighbor connectivity represents a trade-off between two somewhat mutually opposing properties: diversity (for SGD

convergence, also found to be the case in [6, 14] and good neighbor-connectivity (for efficient graph regularization), which usually implies homogeneity. The full batch (i.e., the entire data set) however, has both these properties; perfect neighbor connectivity (since it contains all the points) as well as diversity that mimics the diversity within the complete training data (assuming a large enough, well sampled training set). Indeed, if we are allowed to increase the size of the mini-batches as we please, we could presumably capture a more diverse set of points as well as a significant fraction of their neighbors, but computational and memory constraints prevent us from doing so. Note that the global structure of the affinity graph, owing to its sparsity, consists of a large number of small tightly connected clusters, with relatively few edges between the clusters. Thus a mini-batch that somehow captures this structure, but on a smaller scale, will be expected to have reasonably good connectivity as well as high entropy. We use the following heuristic (described in [13], but reproduced here for clarity) for the construction of improved mini-batches:

1. Given N data points, a batch size B (that represents our memory constraint) and M classes, partition the entire graph into $\frac{NM}{B}$ mini-blocks, where each mini-block is approximately balanced at size B/M.
2. Construct N/B "meta" batches of size B from the mini-blocks as follows:
   (a) For each batch b_i, randomly choose M mini-blocks from the set of $\frac{NM}{B}$ mini-partitions that were created in Step 1.
   (b) Group these M mini-blocks into one larger meta-batch. Since each mini-block is approximately size B/M, our meta-batch will be approximately of size B, satisfying our memory constraint.

At the end of this process we have meta-batches which are of the same size B as the earlier graph-based batches, but which are qualitatively different. Each meta-batch is now composed of many small homogeneous mini-blocks which, due to random sampling, are likely to be of a different class. We omit the proof here due to space constraints, but intuitively we expect that the resulting entropy from grouping together M such randomly chosen mini-blocks (of approximately equal size) to approach the entropy of the training set. The plot in Figure 2a confirms this: as we can see, the mean meta-batch entropy is fairly close to the global entropy, and has a tight distribution around this mean. We also plot the entropy distribution of pure graph-based batches to show the significant level of improvement of batch quality– in terms of entropy – using meta-batches.

With respect to neighbor connectivity, let C_{mini} and C_{meta} denote the random variables that represent the within-batch connectivity of a mini-block and meta-batch respectively. One can show that grouping K mini-blocks to form a meta-batch does not adversely impact the connectivity score, i.e., $\mathbb{E}[C_{meta}] \geq \mathbb{E}[C_{mini}]$. Further, using the Central Limit Theorem, we can show that the variance of c_{meta} is given by $\sigma^2_{c_{meta}} = \frac{1}{K}\sigma^2_{c_{mini}}$. The plots shown in Figure 2b depicting histograms of various batch connectivity scores confirms our analysis. As expected, the green histogram which shows distribution for the meta-batches has approximately the same mean as the mini-partitions (blue histogram) from which it is formed, but a much lower variance.

2.2 Stochastic Neighbor Regularization

Even though a meta-batch constructed using the procedure described in the previous section has much better neighbor-connectivity than a randomly shuffled batch, for a given node, a significant number of neighbors still lie outside the meta-batch. Rregularizing against all neighbors quickly becomes a computational bottlenecl. To preserve efficiency while still regularizing against out-of-batch neighbors, at each step, we *randomly* pick one additional meta-batch and regularize against this neighbor as follows: consider the graph induced by the meta-batches, $\mathcal{G}_M = (V^M, E^M)$, $V^M = \left\{M_1, M_2 ... M_{\lfloor N/B \rfloor}\right\}$ where each M_i is a meta-batch, and edge $e^M_{i,j} \in E^M$ exists between M_i and M_j if there exist some edge $e_{s,t}$ between nodes v_s and v_t in the affinity graph $\mathcal{G}$, such that $v_s \in M_i$ and $v_t \in M_j$. That is, meta-batches are connected if their member nodes are connected in the original affinity graph. Let $\mathcal{C}_{i,j}$ denote the set consisting of all such unique pairs v_s, v_t. Then we can define an edge-weight on each of the edges in E^M as $|\mathcal{C}_{i,j}|$. For a given meta-batch M_j, during each epoch, the probability of picking a neighboring meta-batch M_j is given by

$$p_{i,j} = \frac{|\mathcal{C}_{i,j}|}{\sum_j |\mathcal{C}_{i,j}|}. \tag{6}$$

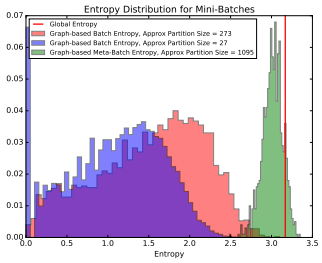
Entropy Distribution for Mini-Batches
Global Entropy
Graph-based Batch Entropy, Approx Partition Size = 273
Graph-based Batch Entropy, Approx Partition Size = 27
Graph-based Meta-Batch Entropy, Approx Partition Size = 1095
Entropy

(a)

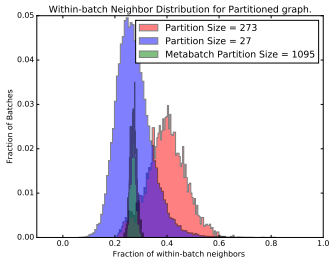
Within-batch Neighbor Distribution for Partitioned graph.
Partition Size = 273
Partition Size = 27
Metabatch Partition Size = 1095
Fraction of Batches
Fraction of within-batch neighbors

(b)

Figure 2: (2a) Entropy distribution of mini-batches (calculated over labels in a mini-batch) for randomly shuffled mini-batches and graph-partitioned minibatches. Shuffled mini-batches have a tight distribution, very close to the entropy of the entire data set (shown as the vertical red line). (2a) Entropy distribution for meta-batches, compared to graph-partitioned batches. (2b) Metabatches preserve the average neighbor connectivity score of the smaller-partitions from which they are formed

Thus, a neighboring batch M_j of batch M_i is more likely to be picked during an epoch if there are a relatively large number of edges between the member nodes that comprise M_i and M_j. Over a large number of epochs, graph regularization is likely to take place against all neighboring batches; this enables labels to propagate via a stochastic diffusion process within the connected components of the affinity graph.

2.3 Decomposing Loss Function for Parallel Training

In Section 1 we noted how our original objective function was not easily decomposable as a sum over data points (or mini-batches) due to the graph regularization term. Using the techniques described above, we can facilitate data parallel training by considering again the objective function in terms of entropy and cross-entropy as given in Equation 3.

$$J_i = H^c(\mathbf{t}_i, \mathbf{p}_i) + \gamma \sum_{j=1}^{n} \omega_{ij} H^c(\mathbf{p}_i, \mathbf{p}_j) - (\kappa + \gamma \sum_{j=1}^{n} \omega_{ij}) H(\mathbf{p}_i) + \lambda \left\| \theta \right\| \tag{7}$$

Now, for a given iteration over meta-batch M_r, and its randomly chosen neighbor meta-batch M_s, the points i n the loss function are simply the labeled points in the concatenated batch $M_c = [M_r, M_s]$ while j is the set of all points in M_c. For a k-worker parallel training scenario, during each iteration, there will be k such meta-batches and the gradients are calculated independently over these batches. Because this is now (approximately[1]) decomposable as a sum over (concatenated) meta-batches, the technique can easily work within a parallel SGD framework. We present the results on a synchronized parallel SGD setup in the next section.

3 Experiments

For the experiments in this paper, we use the TIMIT speech corpus [5] and report the frame-level phone classification accuracy. The training set consists of over 1 million speech frames, each frame being a 351-d vector of cepstral coefficients. The output is mapped to a distribution over 39 classes during scoring. For both the sequential and parallel case, we experiment with label ratios of 2%, 5%, 10%, 30%, 50% and 100% by randomly dropping labels from our training set. For the k-NN graph construction, we set $k = 10$ for all the experiments and use the Scikit machine learning library [11] that constructs the graphs using a fast ball-tree search. After symmetrization, affinities are computed by applying a radial basis function (RBF) kernel, such that each entry w_{ij} in the affinity matrix W, $w_{ij} = e^{-\frac{||x_i - x_j||}{2\sigma^2}}$. For graph partitioning, we use the METIS graph partitioning library [7] that uses a fast recursive multi-way partitioning algorithm to give approximately balanced blocks.

[1]The decomposition is approximate since there will generally be cross-partition edges that we are ignoring in a given iteration

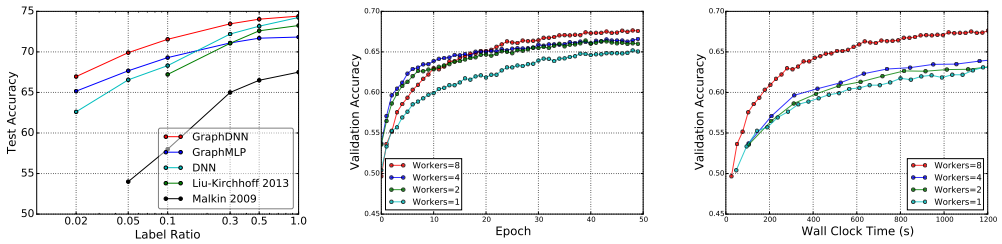

(a) Sequential Case: Final test accuracy

(b) Parallel Case: Validation accuracy vs training epochs

(c) Parallel Case: Validation accuracy vs wallclock time

Figure 3: Sequential and Parallel results for a Graph-Regularized DNN. Fig 3a compares the performance vs. other semi-supervised graph-based schemes in the literature for sequential training. Parallel results are shown in 3b and 3c for a 5% labeled scenario. A more aggressive learning rate for a higher number of workers combined with data parallelism allows us to achieve faster accuracies for a given wallclock time.

Meta-batches are synthesized from the graph partitions as described in the previous sections; both partitioning and meta-batch synthesis are one-time pre-processing steps. We originally implemented all our sequential models using the Theano toolkit [1] and for the parallel case, we implemented our algorithms in the MXNet framework [3] that provides parallel SGD functionality for both synchronous and asynchronous SGD (though we only use the synchronous version in our experiments). For the results reported here we used the AdaGrad [4] variant of gradient descent. We used a DNN with four hidden layers, each 2000 units wide, using Rectified Linear Units [15] as the non-linear activation function, and a softmax output layer. We used dropout while training, reporting the results for the case when dropout probability is 0.2. The model and hyperparameters are all kept unchanged between parallel and sequential versions.

Detailed results for the sequential training case are described in [13] and we reproduce one of the results to illustrate the efficacy of the method (see Figure 3a). In the low labeled scenario, our method significantly outperforms similar graph-based (sequential) SSL methods described in [9, 10]. For the parallel training case, we experimented with 2, 4 and 8 workers, all running on GPU-enabled machines. We keep the batch size the same irrespective of the number of workers, implying that with more workers we average the gradients over a larger number of points. This obviously leads to fewer gradient updates per epoch, compared to the sequential case. However, since gradients are less noisy when averaged over a larger number of training points, in the parallel case we can be a little more aggressive with our learning rate. In our experiments we use a base learning rate of 0.001 and an effective initial learning rate of $0.001k$ where k is the number of parallel workers. After a fixed number of epochs (10 in our case), we reset the learning rate back to 0.001. Thus even the though sequential version has a larger number of gradient updates per epoch, as we can see in Figure 3b, the parallel training run is able to achieve higher validation accuracies per epoch due to a higher learning rate. A more important metric for parallel machine learning in general is the speed of learning – in terms of validation loss – as a function of wall-clock time. We show this in Figure 3c where we see that the eight worker case is able to achieve significantly higher accuracies early on. We should note that synchronizing with the parameter server does introduce an overhead in our experiments – in MXNet, for our implementation, we observe this to be a constant factor of two. That is an individual worker node in the parallel setting was about twice as slow compared to a worker in the sequential version, in terms of sample processing throughput. With a moderate number of workers, we are able to compensate for this overhead, and generally, the method scales well in the data-parallel paradigm, and we expect to see even larger speed ups as we increase the worker count.

4 Conclusion

We presented an efficient, scalable method for parallel graph-based SSL learning that can be applied in general to any parametric learner. The methods presented were heuristically motivated; for our current research we are looking at further analysis, asynchronous versions of SGD and more provably optimal methods for constructing meta-batches.

References

[1] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: a cpu and gpu math expression compiler. In *Proceedings of the Python for Scientific Computing Conference (SciPy)*, volume 4, page 3. Austin, TX, 2010.

[2] Olivier Chapelle, Bernhard Schölkopf, Alexander Zien, et al. *Semi-supervised learning*. MIT press Cambridge, 2006.

[3] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *CoRR*, abs/1512.01274, 2015.

[4] John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12:2121–2159, 2011.

[5] John S Garofolo, Lori F Lamel, William M Fisher, Jonathan G Fiscus, and David S Pallett. Darpa timit acoustic-phonetic continous speech corpus cd-rom. nist speech disc 1-1.1. *NASA STI/Recon Technical Report N*, 93, 1993.

[6] K. Wei, R. Iyer, S. Wang, W. Bai, J. Bilmes. How to intelligently distribute training data to multiple compute nodes: Distributed machine learning via submodular partitioning. In *Neural Information Processing Society (NIPS) Workshop*, Montreal, Canada, December 2015. LearningSys Workshop, `http://learningsys.org`.

[7] George Karypis and Vipin Kumar. Multilevel k-way partitioning scheme for irregular graphs. *Journal of Parallel and Distributed computing*, 48(1):96–129, 1998.

[8] Mu Li, Dave G. Andersen, and Alexander J. Smola. Graph partitioning via parallel submodular approximation to accelerate distributed machine learning. *CoRR*, abs/1505.04636, 2015.

[9] Yuzong Liu and Katrin Kirchhoff. Graph-based semi-supervised learning for phone and segment classification. In *INTERSPEECH*, pages 1840–1843, 2013.

[10] Jonathan Malkin, Amarnag Subramanya, and Jeff A Bilmes. On the semi-supervised learning of multi-layered perceptrons. In *INTERSPEECH*, pages 660–663, 2009.

[11] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *The Journal of Machine Learning Research*, 12:2825–2830, 2011.

[12] Amarnag Subramanya and Jeff A Bilmes. Entropic graph regularization in non-parametric semi-supervised classification. In *Advances in Neural Information Processing Systems*, pages 1803–1811, 2009.

[13] Sunil Thulasidasan and Jeff Bilmes. Semi-supervised phone classification using deep neural networks and stochastic graph-based entropic regularization. In *2016 Workshop on Machine Learning in Speech and Language Processing*, San Francisco, CA, September 2016.

[14] Kai Wei, Rishabh Iyer, Shengjie Wang, Wenruo Bai, and Jeff Bilmes. Mixed robust/average submodular partitioning: Fast algorithms, guarantees, and applications. In *Neural Information Processing Society (NIPS)*, Montreal, Canada, December 2015.

[15] Matthew D Zeiler, Marc'Aurelio Ranzato, Rajat Monga, Min Mao, Kun Yang, Quoc Viet Le, Patrick Nguyen, Alan Senior, Vincent Vanhoucke, Jeffrey Dean, et al. On rectified linear units for speech processing. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 3517–3521. IEEE, 2013.

[16] Xiaojin Zhu, John Lafferty, and Ronald Rosenfeld. *Semi-supervised learning with graphs*. Carnegie Mellon University, Language Technologies Institute, School of Computer Science, 2005.